A Quick Start Guide to Inference Anchoring Theory (IAT)

Centre for Argument Technology

September 25, 2023

Introduction

Inference Anchoring Theory (IAT) [1, 2] provides a theoretical scaffolding to handle dialogue and argument structures, and the relations between them, and has been applied to over 3.5 million words in fifteen languages (available online at `corpora.aifdb.org`). IAT has three types of relations: (i) relations between locutions in a dialogue, called transitions; (ii) relations between content (propositional content of locutions) and (iii) illocutionary connections that link locutions with their content. These relations are collectively known as S-nodes (because they are governed by schemes) and they typically connect propositions known as I-nodes (because they express information). We use IAT in order to represent and understand the arguments people offer, often on public platforms, in order to develop insight into complex debates.

The smallest units of the IAT analysis are *argumentative discourse units* (ADUs). The way they are segmented is described in 1.1. ADUs are typically directly analysed as *locutions*. Locutions typically have speakers (a term we use to encompass utterers in any medium) and typically also have timestamps. The text of locutions should not be reconstructed or changed in any way in contrast to the propositional content of those locutions which be revised to reflect anaphoric resolution and other forms of grammatical repair. Propositional content is (roughly) a sentence that can be assigned a truth value – something that is true or false in the world. So a grammatical sentence such as "Bob likes bananas" can be true or false whereas a noun phrase like "yellow bananas" cannot.

Discrete argumentative function between the propositions is captured by the *propositional relations* of inference, conflict or rephrase described in 2.2. Our goal as analysts is to capture the arguments that were delivered regardless of whether they are good, bad, correct or agreeable to us as analysts. Locutions and propositions are connected via *illocutionary connections*. Some types of illocutionary connections that we often use in the Centre for Argument Technology are described in 3.

Analysis can be conducted using the OVA3 tool `ova.arg.tech`, described in 5.

1 Dialogue structure

1.1 Segmentation and locutions

Segmenting text (or transcribed speech) into units called ADUs is the first step of the IAT analysis.

> An **argumentative discourse unit (ADU)** is any text span which (a) has a **propositional content** anchored in either the locution (ADU) itself or a transition targeting this locution; *and* (b) has **discrete argumentative function**, in that the propositional content stands in relation to one or more other propositions via one or more instances of inference, conflict or rephrase (described in Section 2.2)

Text is segmented into locutions encapsulating ADUs according to the following rules:

1.1.1 **Basics**. Punctuation, delimitation, discourse indicators and other extraneous material that occurs at the boundaries of ADUs are always excluded from the ADU proper. For example, "So, bananas are yellow." should be analysed as a locution, "bananas are yellow" missing out the initial discourse indicator "So", missing out the preceding comma and the terminating full-stop.

1.1.2 **Discourse regulation**. Speakers often work hard to regulate a discourse, explicitly talking about thematic flow (bringing topics back into scope, closing issues down, etc.), turn-taking, and other organisational features of talk. Discourse regulators are almost always excluded from ADUs, whether single lexical terms (*first*, *second*) or much more complex material which is nonetheless entirely regulatory, such as:

> Fiona Bruce: *Let's hear from the person in the blue sweater.*
>
> Audience Member: *Yeah, so it's going back to this point about the council tax being put up by five per cent.*
>
> Fiona Bruce: *Oh hang on, we're talking about knife crime now.*

1.1.3 **Speakers**. Speakers are the utterers of locutions – a term used to encompass those responsible for discourse in any medium. Particularly in analysis of dialogues, though, speakers of particular utterances are identified as part of text of locutions by a convention: "SPEAKER-FORENAME SPEAKER-SURNAME: *ADU*". You should try to go back in the transcript backwards to try to identify the speaker, if it is not immediately clear (because of the use of a pronoun or initials etc). In case it is impossible, include as a name "Unknown" as the first name and "Speaker" as the last name. If you have just one name for the speaker such as The Government, Mr. Mayor, then you can use the article or the title as the first name. If you have only one word such as John or Smith, then you should try to identify the other part of the name.

1.1.4 **Time**. Locutions are typically analysed in temporal order and may also have timestamps.

1.1.5 **Discrete Argumentative Function**. In many cases, a span (any amount of contiguous text) will combine clauses that could be identified as separate ADUs. Examples include conjunctions ("A and B"), conditional clauses ("If A then B"), epistemic modalities ("I think that A") and reported speech ("Bob said that A"). In every case, each span with discrete argumentative function should be analysed separately:

(a) **Conjunctions** are typically analysed as the two constituent conjunct ADUs because the conjunction itself rarely has discrete argumentative function (except in cases such as the natural deduction mechanism of 'and'-introduction)

(b) The analysis of **conditionals** varies: any combination of the three spans ("If A then B", "A", and "B") might have discrete argumentative function depending on the context

(c) **Epistemically qualified statements** are typically analysed as a single segment that drops the epistemic modality – e.g. "I think that A" is typically analysed as just, "A", and "I don't think that A" is typically analysed as just, "not A"

(d) **Reported speech** is almost always analysed into two ADUs: the first corresponding to the complete span and the second to what was reported to have been said

(e) Splitting **utterances such as "Yes, but A"** (or "No, A") have two segments: "Yes" and "A" ("No" and "A" resp.)

(f) **Interposed material** In OVA3, interposed text can present a problem, e.g., "the liquid, because it is so dangerous, is not allowed in the building". In such cases, the text that is interrupted should be identified as one segment, and the interposed text as the other – i.e., in this example, there are two segments: (a) "the liquid, because it is so dangerous, is not allowed in the building"; and (b) "it is so dangerous". The proposition for the first segment can then be edited to remove the interposed text and the second segment needs to be selected again as a new locution. In the case of interposed word order that interacts with reported speech such as "Bob, who is an expert, said A", the same rules are applied, allowing three segments to be identified: "Bob, who is an expert", "A", "Bob said A", (the last one is obtained by selecting the whole span of text and then by editing the proposition node to delete "who is an expert").

1.2 Transitions

Transitions (TA) connect locutions.

1.2.1 **Basics** A transition captures a response or a reply and embodies a functional relationship between predecessor locution and successor locution. A way to check whether there should be a TA connection is the *Wittgenstein test*. TAs capture temporal relation as well, but only to a limited extent - in the sense that the response will always happen later than the locution to which it responds. But the response can refer to something that happened much earlier in the dialogue. Remember that ordering of locutions in discourse is captured by timestamps on locutions, not by transition structure.

> The **Wittgenstein test**. Wittgenstein proposed the notion of a dialogue game to describe how interpersonal discourse is conventionalised. A transition captures that a rule of one of these games has been applied. Think about a game of chess: white swapping king and castle follows a legal move called castling. Black responding by moving a knight diagonally doesn't follow any rule at all. If you're trying to work out if two locutions are connected by a TA, ask: do they follow a rule of how dialogues (of the kind at hand) typically work? Providing an answer to a question is a standard kind of rule of dialogue; whereas suddenly talking about something else does not.

1.2.2 **Adjacency** Transitions often hold between adjacent segments, but not always. It is perfectly common for transitions to skip one or two locutions, and a significant minority capture long-distance relationships in cases where, for example, a claim is returned to and given additional support, or an earlier question is refined, and so on.

1.2.3 **Directionality** Because transitions capture a functional response relation, they very rarely hold in opposition to temporal flow. That is, the directionality of transitions is usually the same as the temporal ordering, where such temporal ordering exists. Because the transition structure can branch, as in a linked argument structure, it is not necessarily possible to reconstruct a complete ordering over all locutions from TA connections alone. To do that, we can use the timestamps on locutions.

1.2.4 **Mirroring propositional relations** If there's a propositional relation on the left side (RAs, CAs and MAs), the TA should typically be between the locutions of the corresponding propositions. The same is done in case of a long-distance relation on the side of propositional relations: you also need to annotate the long-distance TA connection on the dialogical side between the locutions corresponding to these propositions. In this way, you show that the later locution is a response to something which happened some time earlier in the dialogue.

1.2.5 **Non-binary transitions** TAs, like other S-nodes, can have multiple antecedents, but only ever a single consequent (i.e., there can be many arrows in, but just one arrow out of a TA). Thus in some cases, transitions connect more than two locutions. If a given successor locution has two (or more) antecedents, this can be captured with a single TA that has multiple antecedents. In some cases, such a TA can anchor multiple illocutionary forces (just as any TA can). It is bad practice to have multiple TAs that have the same successor locution (because this would necessitate multiple dialogue rules firing simultaneously).

2 Propositional structure

2.1 Propositions

Propositions consist of the contents of individual utterances (locutions). The contents of the propositions needs to be reconstructed as follows:

2.1.1 **Propositional content has to be always a sentence** – it has to have a subject, a verb, a predicate, etc. In other words, you have to reconstruct the missing, implicit material like with anaphoric references.

2.1.2 **Anaphoric references** (words and expressions which refer back to the information that was already given, e.g., pronouns and determiners such as "this", "that", etc.) are typically reconstructed in the text associated with a proposition (i.e., the original text is edited to resolve, e.g., pronouns) where it is possible to do so. **Exception: the pronoun "we" is left as "we", it is not resolved.** Two tricky cases are "They" and "You" which are

resolvable in some cases, but are sometimes intended to mean a general non-specified other, in which case resolution is not possible.

For instance, for the utterance "I didn't say that", you have to reconstruct what she didn't say using the material before this utterance, e.g., if Trump said to Clinton "You said before that taxes should be increased" and she responds "I didn't say that" the propositional content of Clinton's ADU should be: "Hillary Clinton didn't say that taxes should be increased".

2.1.3 In the analysis we distinguish between three types of questions, pure, assertive, and rhetorical. **Assertive and rhetorical questions** need to have their propositional content reconstructed as assertions, e.g., "The weather is good today, isn't it?" will be reconstructed as "The weather is good today". **Pure questions** can be of two main types, information seeking (so, "What time is it?" has content, "The time is *xxx*") and yes/no questions (so, "Is it the case that p?" has the content, "It is or is not the case that p").

2.1.4 **Do as much reconstruction as necessary** so that you end up with a full sentence which will be understandable without any context, i.e., without knowing what has been said before. At the same time, **do as little reconstruction of implicit material as possible** so that you stay close to the original text that you annotate. Think about someone else looking at your analysis without the context of when and who said that.

> The **Standalone Test**. When reconstructing propositions, apply the test: "*Will a third party be able to understand this sentence standalone from any context?*", i.e., is it possible to understand the meaning of the proposition without any additional information. So for example, from "I didn't say that," a reader cannot recognise who the speaker is to whom reference is being made, nor can they know what was(n't) said.

2.2 Propositional relations

Relations between propositional contents are about a speaker's (intended) use of linguistic material. We distinguish three types of propositional relations, Inference (RA), Conflict (CA), and Rephrase (MA). All RAs, CAs, and MAs must be anchored through Illocutionary Connections (ICs) in Default Transitions (TAs).

It is important that as analysts, we allow arguers to express not just good arguments, but poor, weak, incoherent and fallacious arguments. We are asking ourselves if a speaker intended the content of her utterance to be understood to be related to previous material in a given way. We are never trying to analyse the ways in which propositions *might* be connected. It is tempting to see two propositions that naturally look like a premise and a conclusion and want to connect them. But we are not capturing how the (logical) world is, nor how we as analysts would like to argue. We are aiming to capture the arguments that were made by the original interlocutors.

2.2.1 **Inference (Support, RA)** Holds between two propositions when one or more propositions (premises) are used in order to provide reasons to accept another proposition (conclusion). Support may be of a specific kind, depending on the theoretical context an analyst is working in – Modus Ponens ("P implies Q. P is true. Therefore Q must also be true."), Argument from Expert Opinion, and (the prima facie reasoning from) Perception are all examples of such kinds. If a support relation is not associated with a specific scheme, it defaults to 'Default Inference.' Any given support is thus an application of a rule of inference – a **R**ule **A**pplication, hence RA.

There are several patterns of inference structure:

(a) **Serial inferences** Serial inferences occur when there is an inference relation from a first proposition to a second proposition, and another inference relation from the second proposition to a third proposition. Serial inferences necessarily involve multiple RAs.

(b) **Convergent inferences** Convergent inferences occur if there is an inference relation from a first proposition to a second proposition, and an independent inference relation from a third proposition again to the same second proposition. Each inference relation is a separate RA.

(c) **Linked inferences** Linked inferences occur if there is an inference relation from a first proposition together with a second proposition to a third proposition. This happens when the first or second proposition alone cannot support the third proposition. Both premises are antecedent to the same RA.

(d) **Divergent inferences** Divergent inferences occur if there is one inference relation from a first proposition to a second proposition, and an independent inference relation from the same first proposition to a third proposition. Each inference relation is a separate RA.

2.2.2 **Conflict (Attack, CA)** Holds between two propositions when one proposition is used in order to provide an incompatible alternative to another proposition. In other words, one can ask the following question in order to identify conflict: Is it possible for an individual to commit to both propositions and not be accused of being incoherent? If the answer is no, then we can assume that there's conflict between the propositions. As with inference, conflict may also be of a given kind (e.g., Conflict from Bias, Conflict from Propositional Negation) and similarly defaults to 'Default Conflict.' [1] In contrast to inference, conflict is always structurally the same – it only ever has one incoming and one outgoing edge (the one exception is the theoretically possible situation in which a conflict is undercut *q.v.*). There is therefore no interpretation of a "linked conflict" with multiple antecedents. In situations where there are two propositions used together to form a conflict relation with a third, the best analysis is usually to introduce an implicit proposition that is the conclusion of an inference from the first two, or, alternatively, simply to analyse the two propositions as one if they together have discrete argumentative function.

Finally, notice that conflict is a broader concept than logical contradiction, i.e., "This dress is green" can be attacked not only by "This dress is not green", but also by "This dress is blue".

There are several patterns of conflict structures:

(a) **Rebutting conflict** If a conflict relation targets a proposition that forms the conclusion of another argument, then the conflict is rebutting.

(b) **Undermining conflict** If a conflict relation targets the premise of another argument, then the conflict is undermining.

(c) **Undercutting conflict** If a conflict relation targets the inference relation between a set of propositions and another proposition, then the conflict is undercutting.

2.2.3 **Rephrase (MA)** Holds between two propositions when one proposition is used to paraphrase, restate or reformulate another proposition. Rephrase involves different propositions connected through a variety of different relations, such as specialisation, generalisation, instantiation, etc. Question answering often involves rephrasing because the propositional content of a question is stereotypically instantiated, resolved, or refined by its answer.

> **Rephrasing is not repeating**: repetition involves multiple utterances with the *same* (i.e., just a single) propositional content. If the second locution repeats the content of the first locution, then the content of the second locution is the same as the content of the first one. In implementation, identity conditions for propositions are currently effectively string matching (hence the need for anaphoric and deictic reconstruction). So a repetition has quite literally the same propositional content. Notice that the locution may differ slightly and still have the same – and therefore repeated – propositional content.

In contrast to inference (RA), but like conflict (CA), rephrase (MA) structures are almost always the same – they have only one incoming and one outgoing edge (linked conflicts and linked rephrases **are not possible**.) If you have two propositions rephrased as one, this is best analysed as an inference (in natural deduction, this is inference is referred to as 'And-introduction').

[1] Note that conflict need not be symmetric. Some kinds (such as Conflict from Propositional Negation) typically are symmetric, which should be captured with two distinct Conflict relations, one in each direction. Such symmetric attacks are very rare in analysis of textual material.

3 Illocutionary structure

Illocutionary connections link both locutions and transitions between locutions (TAs) to both propositions and relations between propositions – RA, CA and MA (that is, they link the right hand side of the IAT diagram with its left hand side). Locutions and transitions are said to **anchor** illocutionary connections (and sometimes also to anchor their propositional contents). The illocutionary connections are applications of illocutionary force introduced by Speech Act Theory. Illocutionary connections represented in YA nodes are governed by schemes, just like RA, CA and MA; each YA scheme captures a single illocutionary force named by convention with a gerund. If am illocutionary connection is not associated with a specific scheme, it defaults to 'Default Illocuting'.

Each locution or transition will often anchor a single YA, but may sometimes anchor more than one. There is no single prescribed set of illocutionary schemes (any more than there is are prescribed sets of inference, conflict or transition schemes). Illocutionary schemes suitable for negotiation (that might involve, e.g., *offer* and *reject*) might be different from those involved in a court (where *testify* and *object* might be more appropriate).
That said, many domains share a number of schemes, and we adopt a restricted set in much of our annotation work:

3.1 **Asserting (A)** The speaker S is *asserting* p to communicate their opinion on p. It does not imply that S really believes p: it is rather a public declaration to which the speaker can be held.

3.2 **Questioning (Q)** S is *questioning* p when S formulates p as an interrogative sentence with the form of a Wh-question or a Yes/No or closed-list question. We distinguish three categories of questioning: Pure Questioning (PQ), Assertive Questioning (AQ) and Rhetorical Questioning (RQ). AQ and RQ, in contrast to PQ, carry some degree of assertive force.

- **Pure Questioning (PQ):** S is asking for the hearer H's opinion on p: whether H believes p, or not, or has no opinion on p. The propositional content of PQ is treated as underspecified – as a disjunction for a yes/no or closed-list question (i.e., "Is it the case that p?" has the content, "It is or is not the case that p") or as a lambda sentence for a Wh-question: (i.e., "What time is it?" has content, "The time is xxx").
- **Assertive Questioning (AQ)** For AQ, S not only seeks H's opinion on p, but also indirectly publicly declares their own opinion on p. It is typically linguistically strongly signalled by cues such as "Isn't it the case that...", "Can we agree that...", "Doesn't...".
- **Rhetorical Questioning (RQ)** Finally for RQ, S is grammatically stating a question, but in fact is only conveying that they do (or do not) believe p. A good test for deciding between RQ and AQ is to check whether it is discursively reasonable for H to reply to the question, e.g., whether their response "Yes" to the question "Does the pope wear a funny hat?" would be treated as irrational (or humorous or naive) discursive behaviour.

3.3 **Challenging (Ch)** When S is *challenging* p, S declares that they are seeking (asking about) the grounds for H's opinion on p. Challenges are a dialogical mechanism for triggering argumentation.

3.4 **Agreeing (Agr)**Agreeing is used for expressing a positive reaction, i.e., when the speaker S declares that they share the opinion of the interlocutor. This can take the basic form of signalling such as "Yes", "Indeed", "Most definitely", "Sure", but equally may be a complete sentence. Note that the utterance, "Yes" may not on its own anchor agreement. If it is as a response to a question (AQ, PQ, RQ), the agreement is anchored in the transition between the question and the response. Agreeing typically takes as content a proposition that was anchored in an earlier locution.

3.5 **Disagreeing (Disagr)**Disagreeing is used for expressing a negative reaction, i.e., when S declares they do not to share the interlocutor's opinion. This can take the form of utterances which have similar meaning to "No" (e.g. "I'm not saying that", "Actually, that's not correct", "Definitely not", "No, it's not") or it can be an utterance with a complete propositional content. One might expect that disagreeing should have a structure that mirrors agreeing but in fact it is very different. Where agreeing has as its content something previously introduced, disagreement requires the introduction of new material, *viz.* the counterposition. The locution will usually anchor the counterposition through a YA of asserting (or similar). The content

of disagreeing YA itself is the CA between the original position and the new counterposition, and is usually anchored in the TA between the locution of the position and the locution of the counterposition.

3.6 **Restating (Rest)** *Restating* is used for expressing the relation of rephrase between propositional contents, i.e., it anchors the MA between two propositions in the TA between the two corresponding locutions. The rephrase can take the form of an utterance that slightly modifies the original content of the locution being restated.

3.7 **Arguing (Arg)** S is *arguing* when they defend a standpoint. It is signalled by linguistic cues such as "because" and "therefore"; however, these indicators occur infrequently in spoken natural language. Arguing takes as a content a relation of inference (i.e., an RA). In other words, the *inference relation* between two propositional contents – a premise and a conclusion (the left hand side of the IAT diagram) is anchored in the TA between two matching locutions (the right hand side of the diagram) by means of an illocutionary force called *Arguing*.

3.8 **Default Illocuting** If an illocutionary connection does not match the guidelines for any IC described above, then it can be labelled *Default Illocuting*. This illocutionary connection is also currently used to connect an MA to its TA anchor when the MA is being used to answer a question.

One common illocutionary structure concerns **reported speech**. Such reported speech (whether direct or indirect) is handled using a stereotypical pattern of locutions and YAs whereby the locution expressing the reported speech anchors via a YA (typically of asserting) the speech that is being reported. This too is a locution and in its turn anchors a YA through to its own propositional content. This structure reminds us that although we may talk of left- and right-hand sides of an analysis, in fact locutions are just a special type of proposition, so analyses can become much more layered and interconnected.

Finally, all RAs, CAs and MAs must be anchored through YAs in locution and TA structure.

4 Examples

4.1 Segmentation

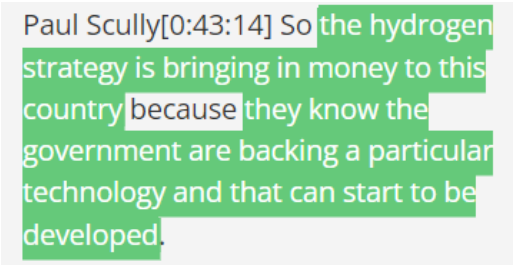

(a) Example with 'so' and 'because'

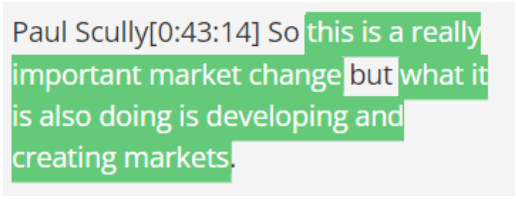

(b) Example with 'so' and 'but'

Figure 1: Conjunctions

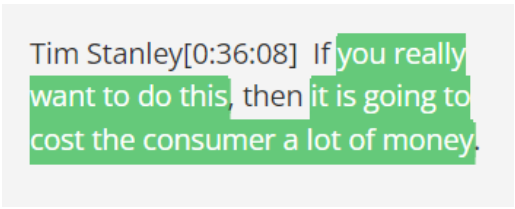

Figure 2: Conditionals

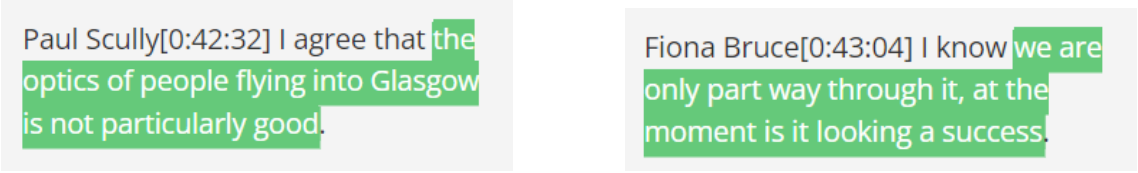

(a) Example with 'I agree'

(b) Example with 'I know'

Figure 3: Epistemic modalities

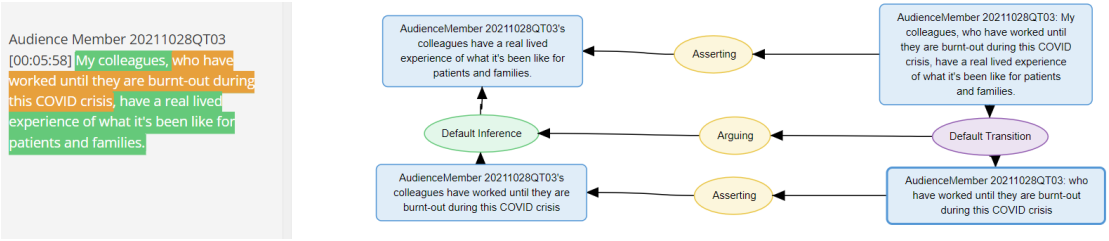

Figure 4: Interposed material

4.2 Reconstruction

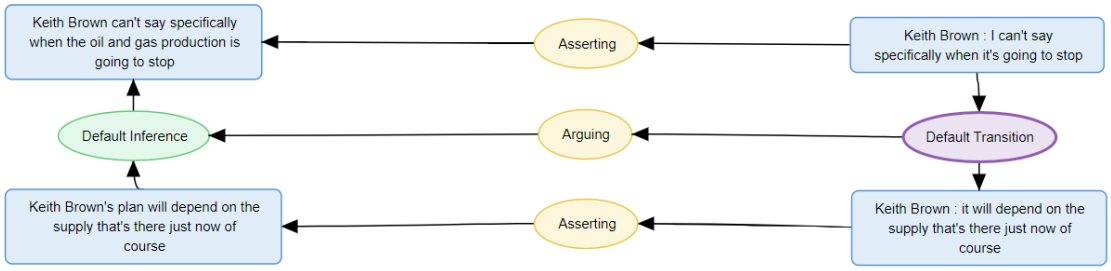

Figure 5: Reconstruction (pronouns and anaphoric references)

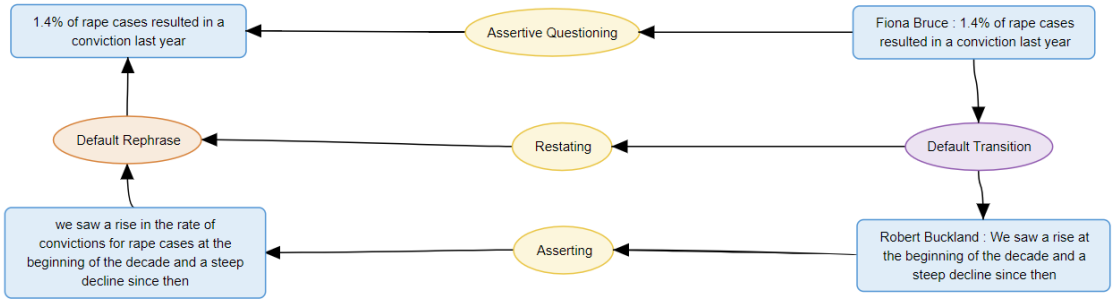

Figure 6: Reconstruction (anaphoric references)

4.3 TA structures

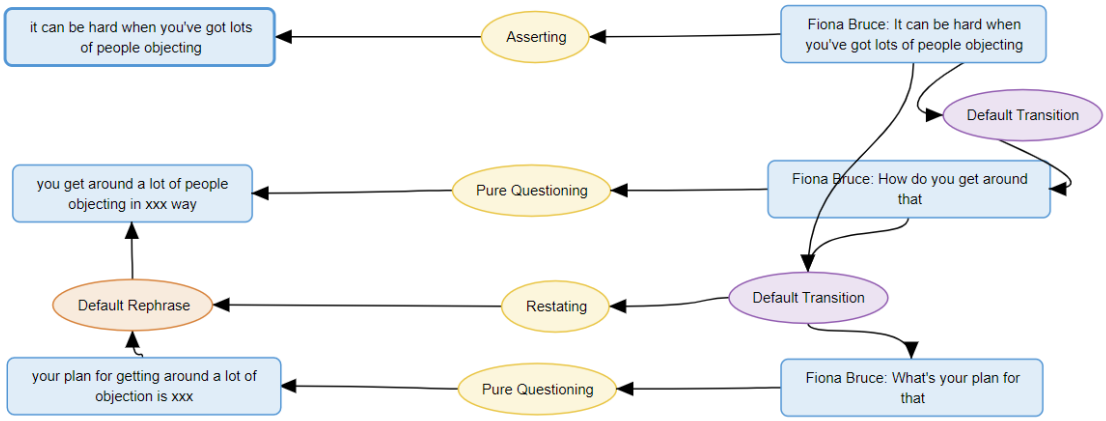

Figure 7: TA structure with two antecedents for one TA

4.4 RA structures

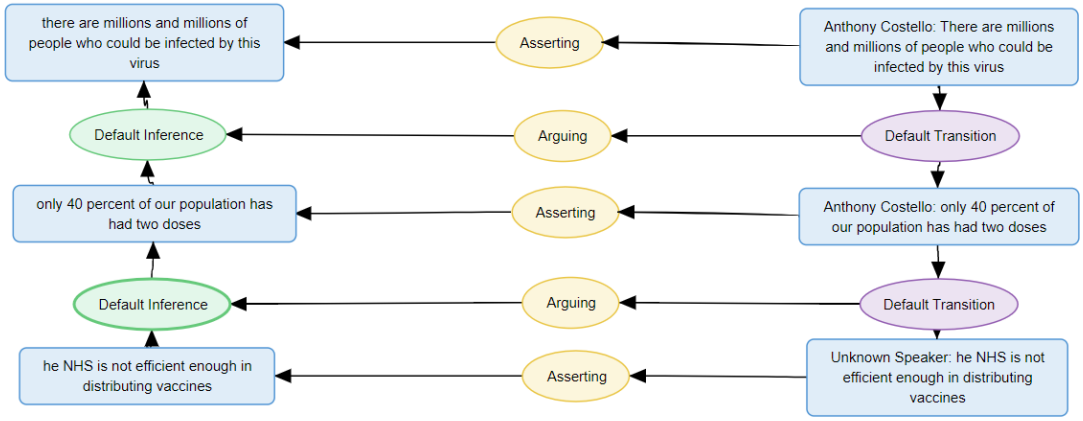

Figure 8: Serial argument

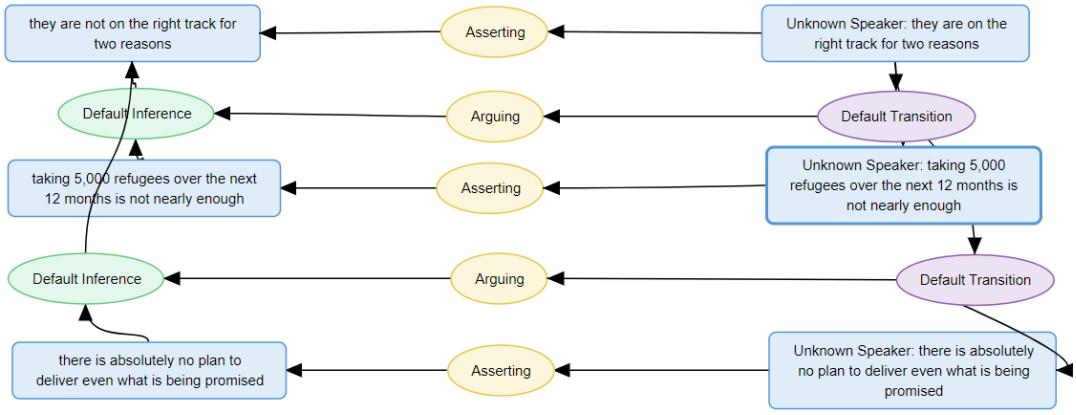

Figure 9: Convergent argument

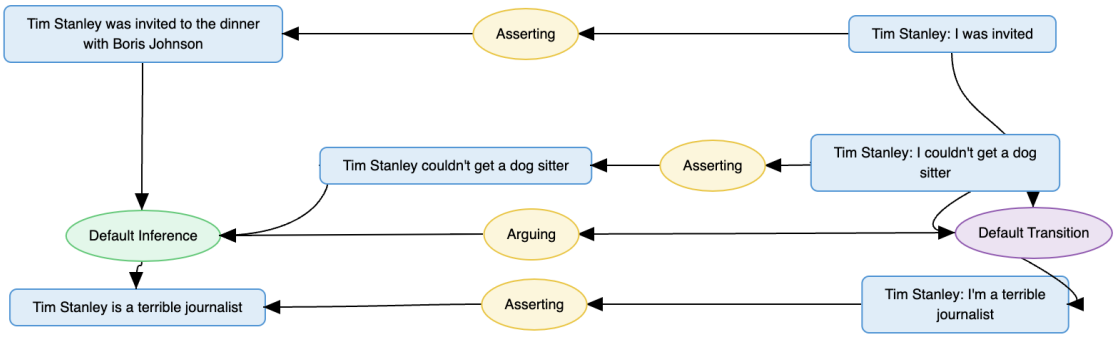

Figure 10: Linked argument

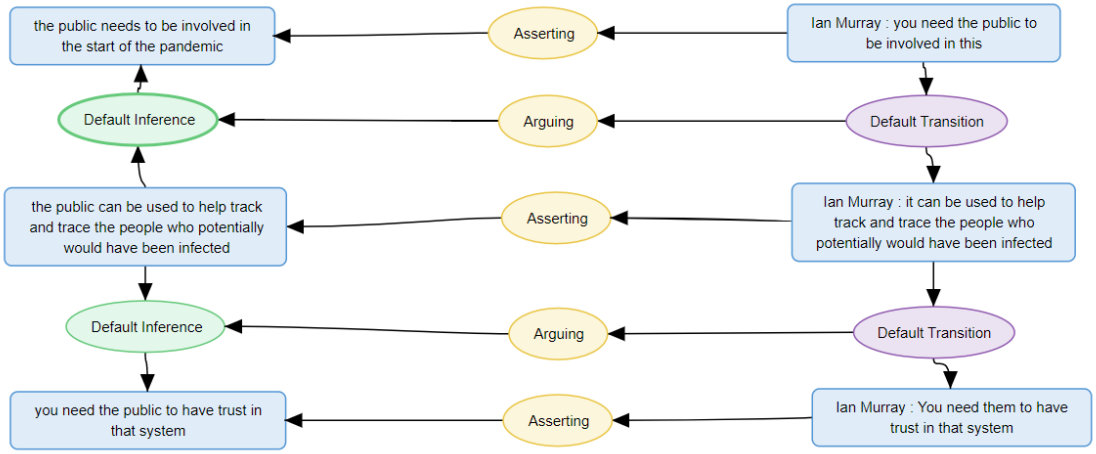

Figure 11: Divergent argument

4.5 CA structures

4.5.1 Rebutting

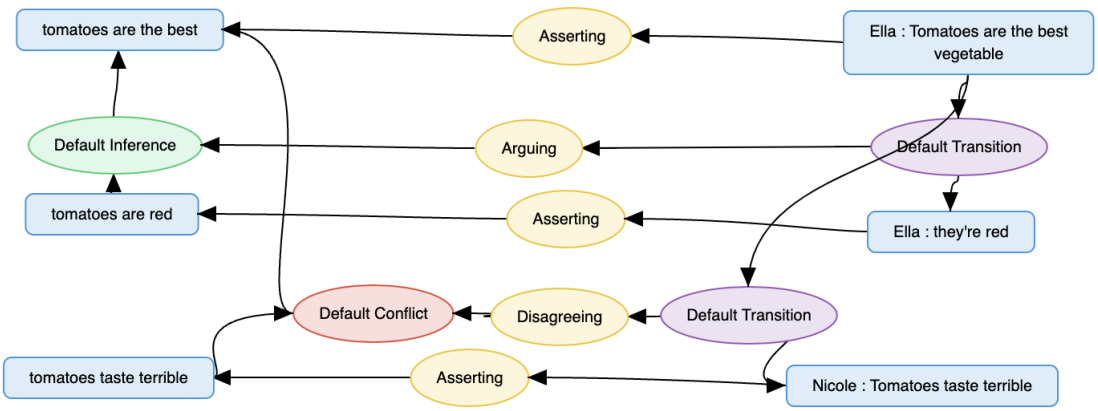

4.5.2 Undermining

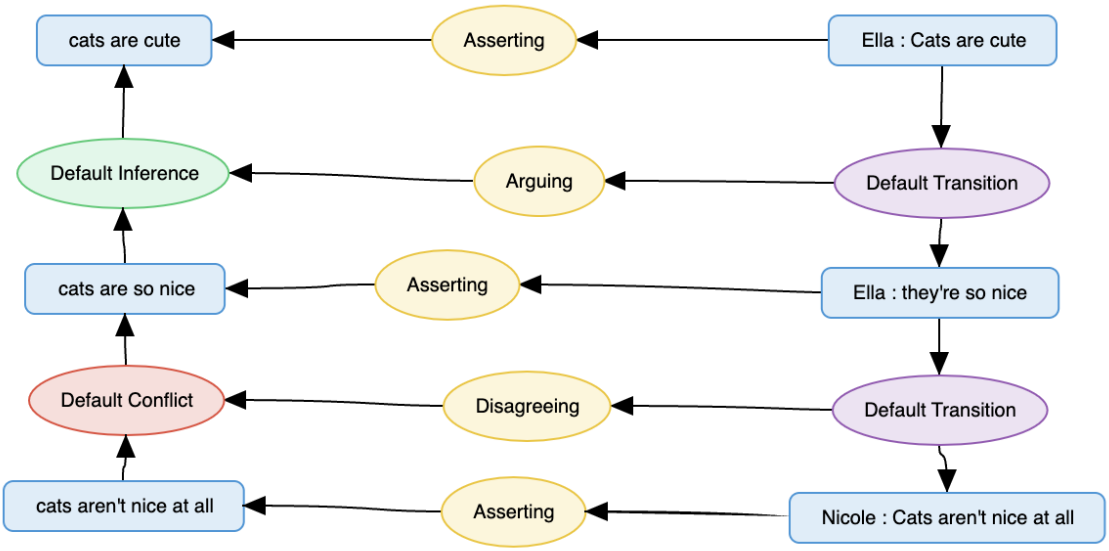

4.5.3 Undercutting

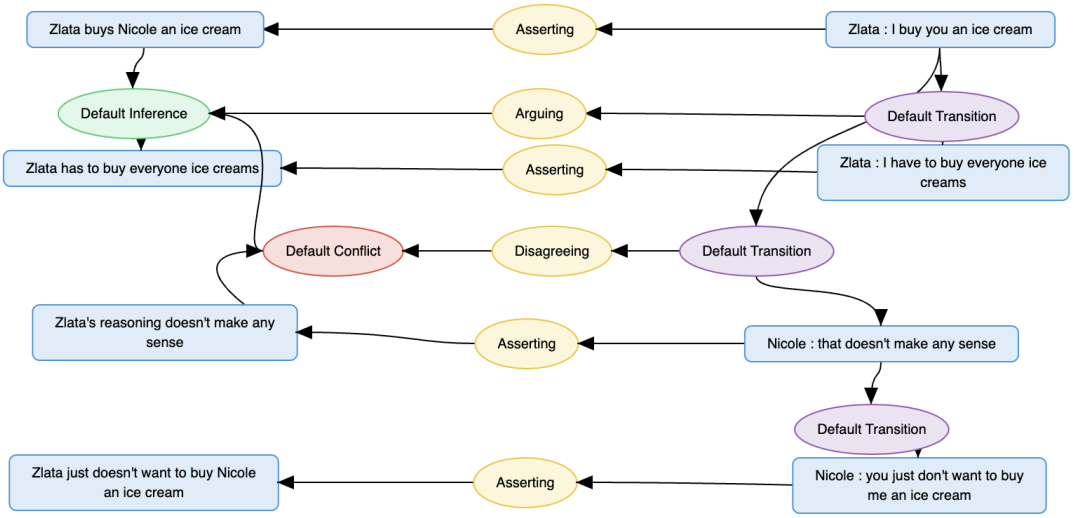

4.6 MA structures

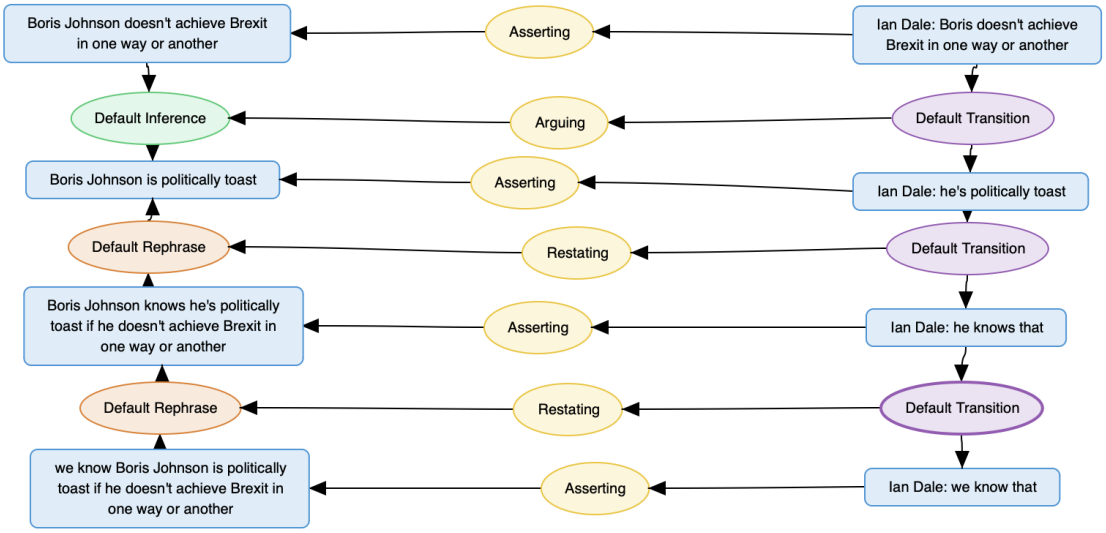

4.7 Questioning

4.7.1 Pure Questioning

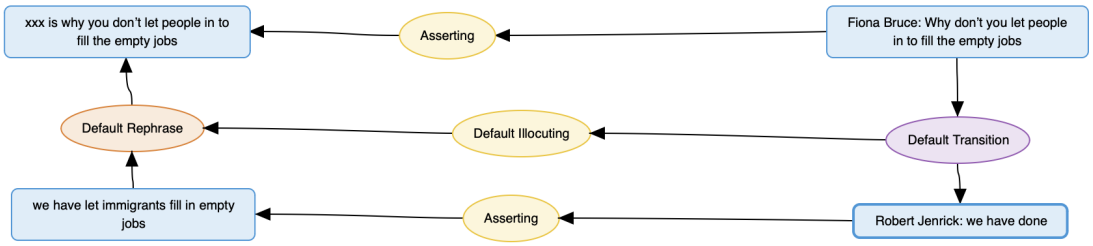

Figure 12: xxx type

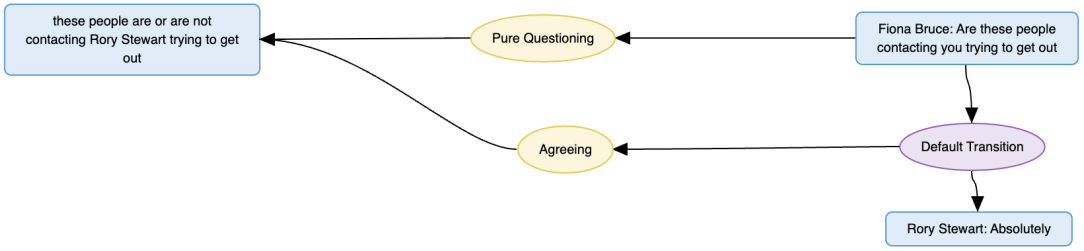

Figure 13: is-or-is-not type

4.7.2 Assertive Questioning

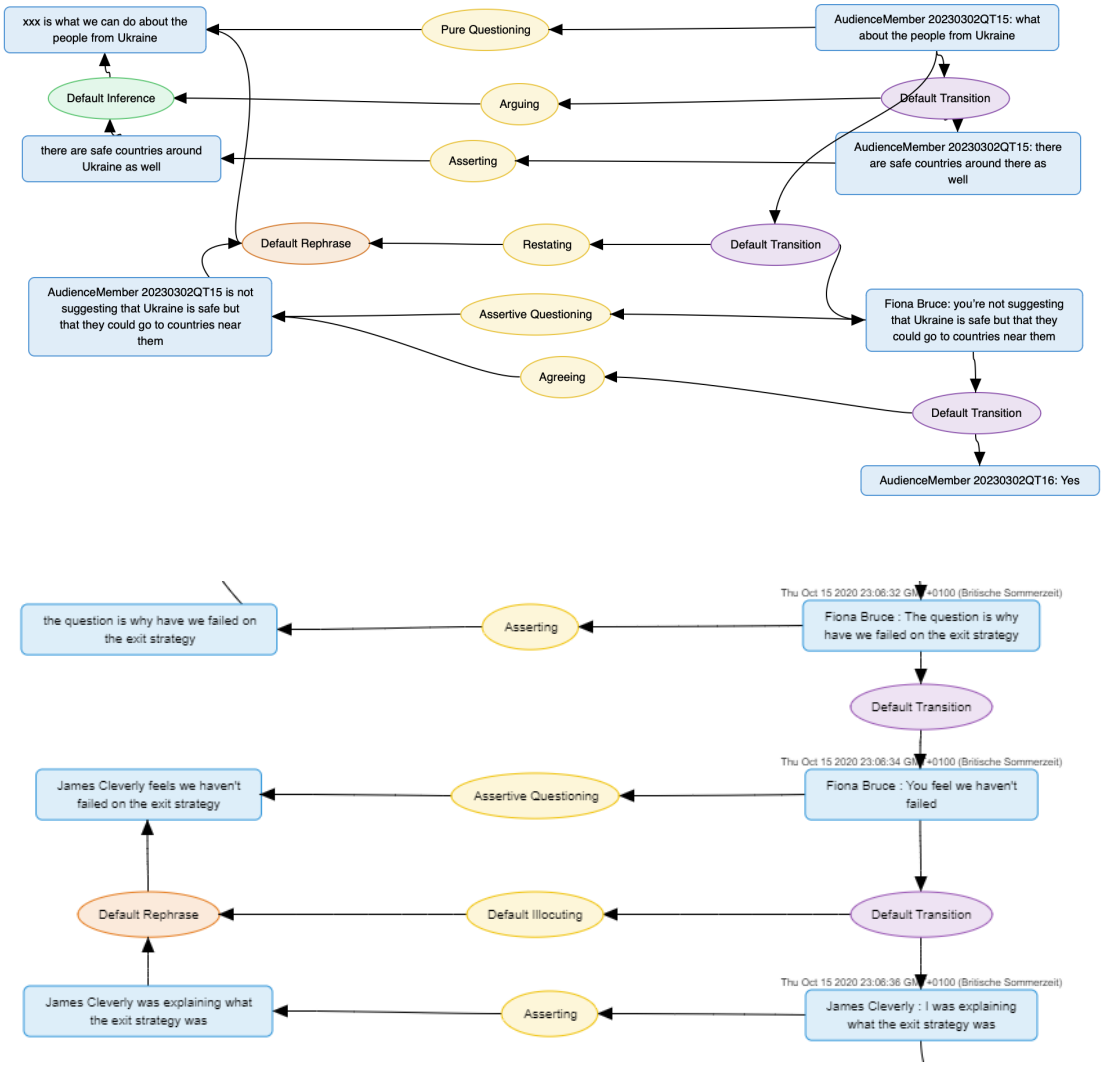

Figure 14: Assertive question example from AIFdb 23809

4.7.3 Rhetorical Questioning

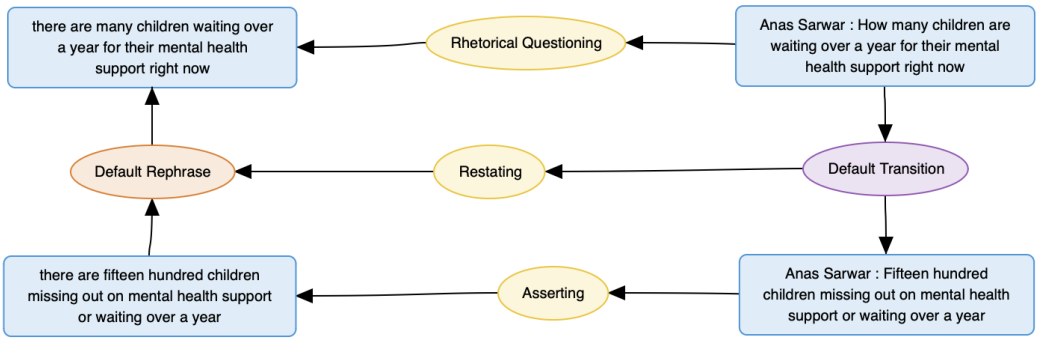

4.8 Challenging

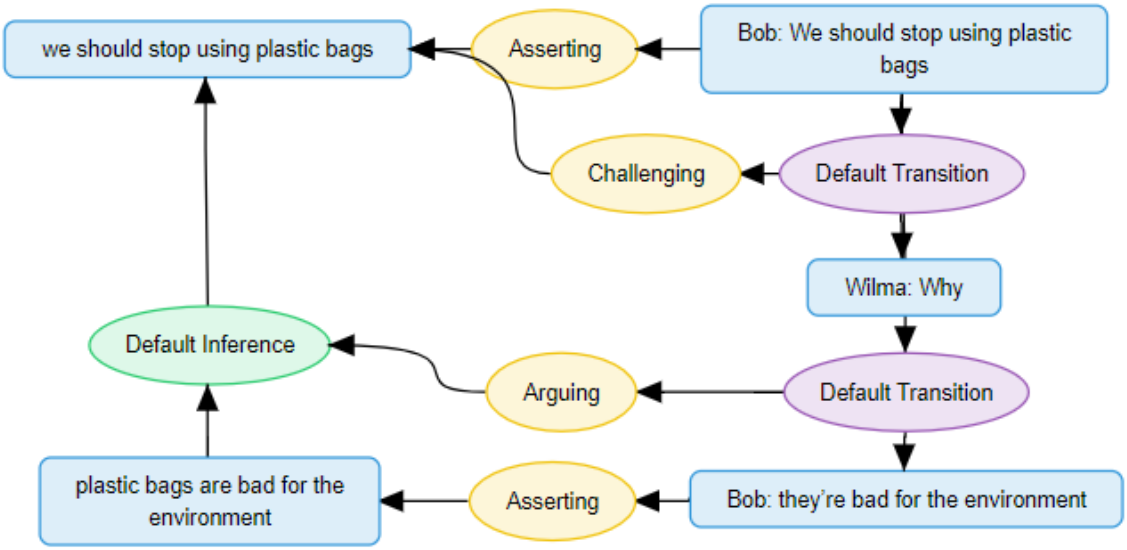

Figure 15: Requesting a reason

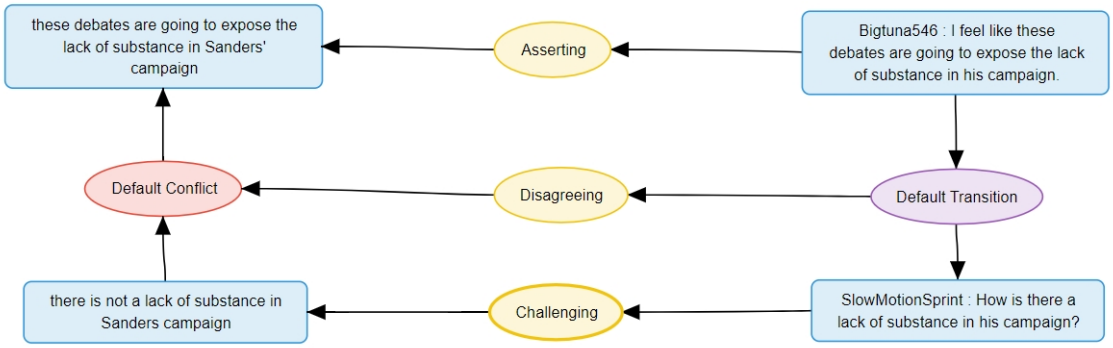

Figure 16: Questioning reasoning example from AIFdb 27381

4.9 Agreeing and Disagreeing

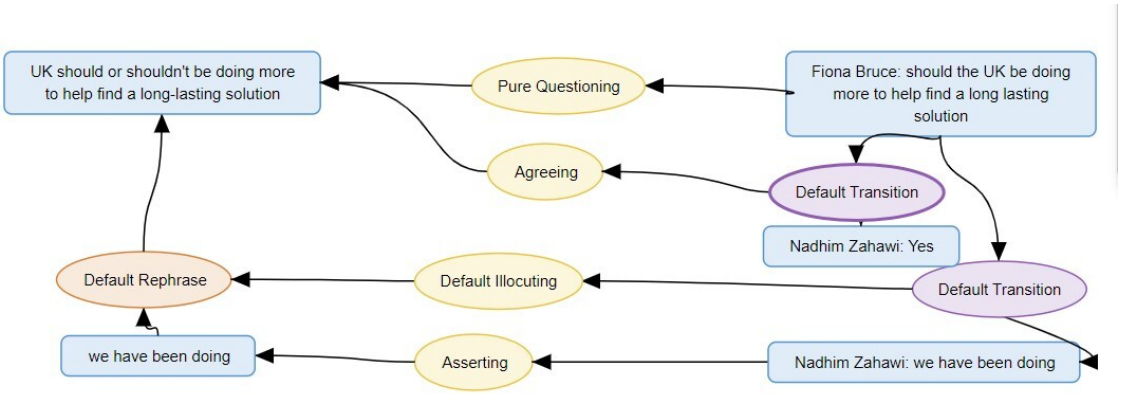

Figure 17: Agreeing with a pure question

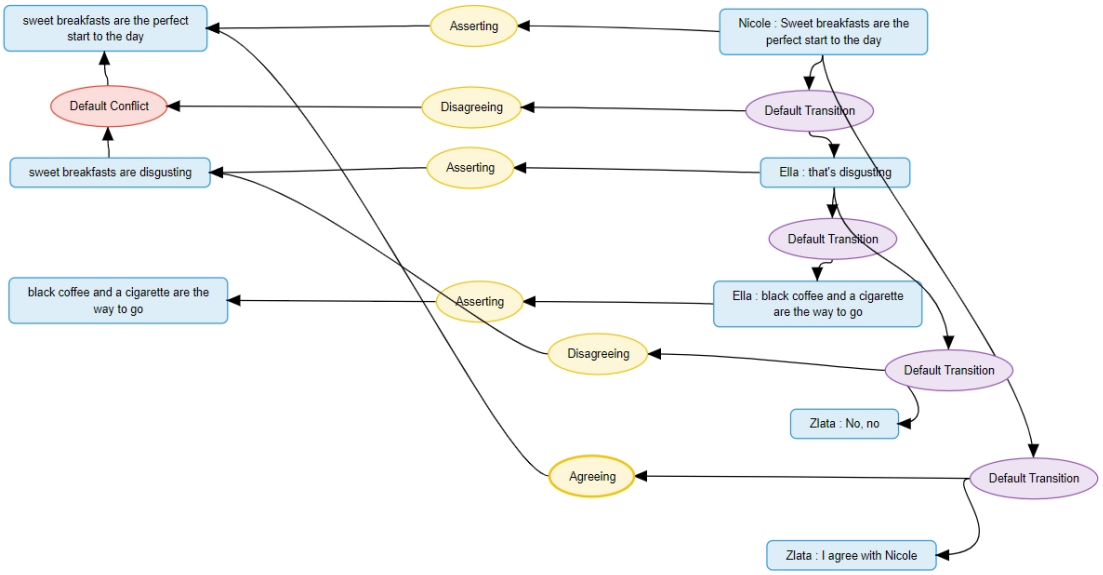

Figure 18: Agreeing and disagreeing example from AIFdb 26941

4.10 Reported Speech

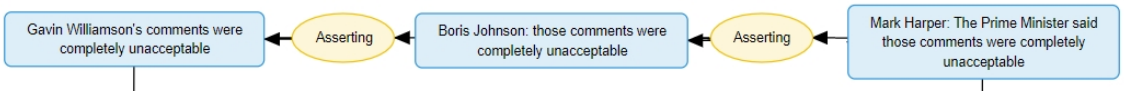

Figure 19: Simple reported speech example from AIFdb 27245

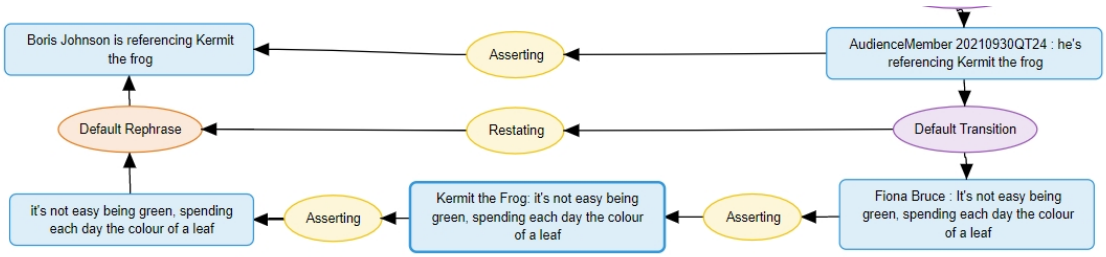

Figure 20: Reported speech of a quote example from AIFdb 27377

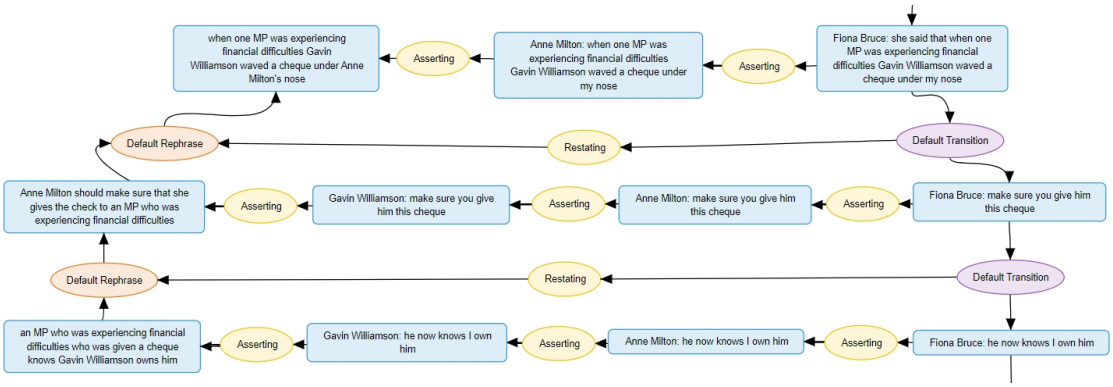

Figure 21: Example of one speaker reporting what a person said that another person said from AIFdb 27245

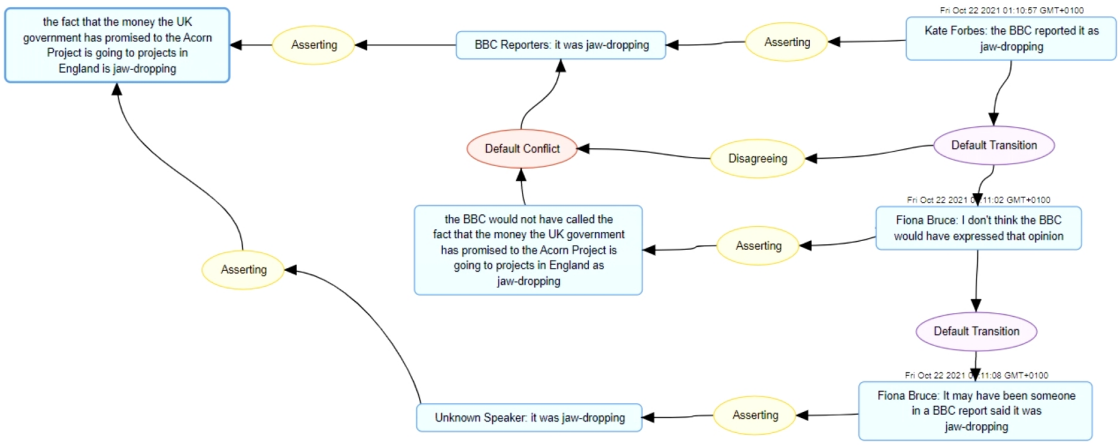

Figure 22: Example of disagreeing with who has been reported as saying something

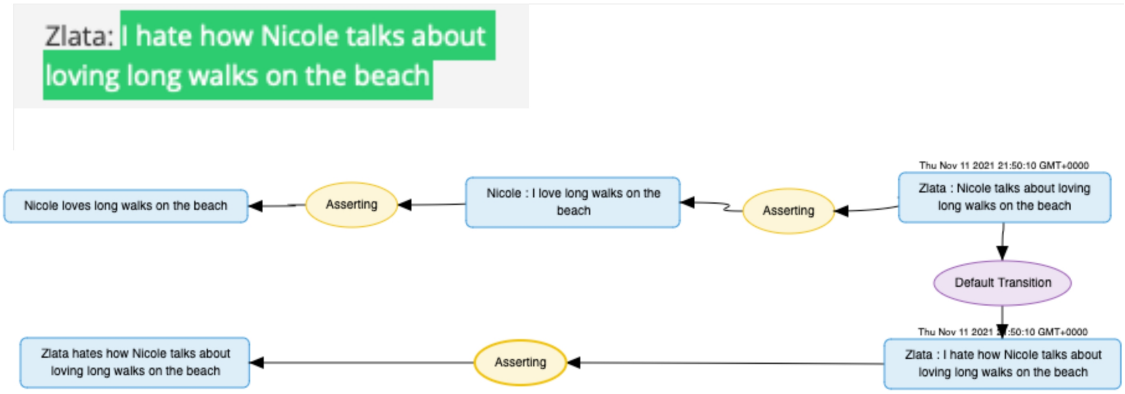

Figure 23: Reported speech as interposed material

5 Annotation software

The OVA3 argument diagramming tool

OVA3 is an online tool for argument analysis facilitating the representation of the structure of argumentative discourse. You can start using OVA3 freely at the website `ova.arg.tech`. A manual for using OVA3 is available at `arg.tech/~chris/OVA3UserGuide.pdf`.

OVA3: Maps' layout The maps should be clearly laid out, without too many overlapping nodes and edges. Otherwise, it is difficult to see, e.g. what is the conclusion of an inference if two arrows are laying on each other and leading to two different nodes.

OVA3: Associating speakers with locutions You can associate a locution with either a pre-existing speaker or a new speaker – in the latter case you can type in the forename and surname (with usual capitalisation) which OVA will then prepend with a colon to the locution content. Once you've added a new speaker they'll be available for subsequent selection.

The AIFdb and AIFdb Corpora repositories

Analyses produced with OVA3 can be saved as 'argument maps' in AIFdb, an online searchable repository of analysed arguments freely available at `aifdb.org`. The argument maps can be collected in corpora at `corpora.aifdb.org`.

List of Figures

References

[1] Katarzyna Budzynska, Mathilde Janier, Chris Reed, and Patrick Saint Dizier. Theoretical foundations for illocutionary structure parsing. *Argument & Computation*, 7(1):91–108, 2016.

[2] Mathilde Janier, Mark Snaith, Katarzyna Budzynska, John Lawrence, and Chris Reed. A system for dispute mediation: The mediation dialogue game. In *The 6th International Conference on Computational Models of Argument*, pages 351–358. IOS Press, 2016.